IBM @server® Cluster 1350

IBM @server Cluster 1350

Highlights

- ***Helps reduce time and resources necessary to deploy Linux® clusters***
- ***Provides single point-of-control to simplify management and enhance cluster availability***
- ***Offers a highly scalable solution for high-performance or commercial applications and server consolidation***

Outstanding price/performance

Linux clustering offers significant price/performance advantages for many high-performance workloads by harnessing the advantages of low cost servers plus innovative, easily available open source software.

Today, many businesses are building their own Linux clusters using commodity hardware, standard interconnects and networking technology, open source software, and in-house or third-party applications. Despite the advantages offered by these systems, the expense and complexity of assembling, integrating, testing and managing these clusters from disparate, piece-part components often outweighs any benefits gained.

IBM has designed the IBM @server® Cluster 1350 to help address these challenges. Now clients can benefit from IBM's extensive experience with clustered UNIX® and Linux computers to help minimize this complexity and risk. Using advanced Intel® Xeon™, AMD Opteron™, and IBM PowerPC® and POWER5™ processor-based server nodes, proven cluster management software and optional high-speed interconnects, the Cluster 1350 offers the best of IBM and third-party technology. As a result, clients can speed up installation of a Linux cluster, simplify its management, and reduce mean time to payback.

The Cluster 1350 is designed to be an ideal solution for a broad range of application environments, including industrial design and manufacturing, financial services, life sciences, government and education. These environments typically require excellent price/performance for handling high performance computing (HPC) workloads. It is also an excellent choice for applications that require horizontal scaling capabilities, such as Web serving and collaboration.

A comprehensive solution

The Cluster 1350 is a flexible, integrated offering that can help optimize time-to-production by reducing the time and resources required to configure, assemble, integrate, test and manage a Linux cluster. IBM consultants can help configure a complete cluster solution, including all the necessary hardware, designed for the client's specific needs. When combined with the client's choice of the Linux operating system[1] and IBM Cluster Systems Management (CSM) for Linux software, the Cluster 1350 provides an easy-to-configure platform to help expedite rapid deployment of cluster applications.

Software installation services are available from IBM Global Services or IBM Business Partners. These services may include planning, installation and configuration of the Linux operating system as well as additional cluster software. For even higher levels of support, clients may take advantage of the optional SupportLine for Linux Clusters. These SupportLine services cover all cluster hardware components, the Linux operating system, and CSM for Linux and General Parallel File System (GPFS) for Linux software.

To further simplify the deployment effort, IBM offers project management support to coordinate all aspects of delivery and installation, including hardware and software set-up services. And, attractive financing and leasing terms are also available through IBM Global Finance.

High-performance cluster management

IBM offers CSM for Linux, an advanced cluster management software tool that allows single point-of-control management for a cluster of supported Xeon, Opteron, PowerPC and POWER5 processor-based servers. This tool helps simplify the management of the cluster and easily scales with the cluster to improve the efficiency of the system administrator.

CSM allows the administrator to monitor both hardware and software events from a single local or remote location, and can generate automatic notification and recovery actions as defined by the administrator. Furthermore, CSM provides node grouping capability, a convenient and essential tool to allow an administrator to apply rules and commands to an individual node, groups of nodes, or the entire cluster. These cluster control event monitoring and automated recovery features help enhance the efficiency of the administrator and the reliability of the cluster by enabling highly effective node control and rapid problem detection and resolution.

In addition to single point-of-control management for all Cluster 1350 server nodes running Linux, CSM is also available for the AIX 5L™ operating system and allows non-switched clusters of IBM System p5™, **@server** p5 and pSeries® nodes running AIX 5L or Linux to be integrated into the same cluster with Xeon Opteron, PowerPC and POWER5 processor-based server nodes. IBM xSeries® and BladeCenter® nodes can also be monitored and controlled using IBM's management software for servers, IBM Director 5.1.

CSM contains several components designed to help simplify the management and increase the reliability and availability of Linux clusters:

- ***Distributed management server:*** *Provides a persistent repository of information about each node in the cluster, and maintains the status of each node.*

- ***Event response resource manager:*** *Provides the ability to run a rich set of predefined commands or scripts in response to user-defined events. Cluster resources that can be monitored include nodes, adapters, file systems and processes.*

- ***Remote hardware control:*** *Utilizes the integrated systems management processor in Cluster 1350 nodes. This feature enables the administrator to remotely reset or power the node on or off.*

- ***Configuration file management:*** *Provides a repository for files that are common between the nodes. CSM synchronizes changes to configuration files across the cluster.*

- ***Distributed shell:*** *Allows commands or scripts to be run remotely on all nodes in the cluster with options for combining output from multiple nodes. Distributed command execution manager is an optional graphical user interface that integrates with the distributed shell, allowing easier management of nodes and node groups.*

Flexibility and advanced server technology
The Cluster 1350 allows clients to choose the right advanced server technology to meet the specific application needs of their environment by offering configurations built on Xeon, Opteron, PowerPC and POWER5 processor-based servers.

The Intel Xeon servers are xSeries 336 and 346 and BladeCenter HS20. Their design leverages IBM Enterprise X-Architecture™, which combines key availability features inspired by IBM zSeries® servers with the advanced scalability features of pSeries systems. All of these servers offer affordable enterprise power and control for 32- or 64-bit processor-intensive applications in an Intel Architecture (IA) environment. In addition, the BladeCenter HS20 provides additional benefits for clients who value the combination of performance, density, integration and investment protection offered by its revolutionary architecture that delivers up to double the density available in 1U (1 EIA Unit) servers.

The Cluster 1350 broadens the choice of superior server technologies with the IBM @server 326m and AMD Opteron LS20 for BladeCenter based on the AMD Opteron processor. Both of these servers incorporate the latest single- and dual-core processors from AMD featuring the unique HyperTransport technology and integrated memory controller to help minimize bottlenecks related to memory access and I/O bandwidth. These servers are especially well suited to computationally intense environments that have previously been constrained by processor-memory bandwidth. In addition, the innovative design of BladeCenter combined with the new LS20 compute node is an ideal choice for clients with memory-intensive applications who value enhanced processor-memory performance and compute density, integration and investment protection in their HPC platforms.

For those clients who prefer a native 64-bit computing platform, the BladeCenter JS20 blade server further extends the Cluster 1350 portfolio by delivering the price/performance capabilities of the IBM PowerPC 64-bit processor in a blade architecture. The PowerPC architecture allows the same operations to be performed on different data sets in parallel, and is especially beneficial to clients who have applications that can exploit the unique capabilities of VMX (Vector Multimedia eXtension). This processor technology combined with IBM's innovative BladeCenter architecture offers clients the opportunity to reach new levels of peak computing capacity in a single rack. And now, clients can also choose the IBM @server OpenPower™ 710 and 720 as the management and/or storage nodes for their JS20-based Cluster 1350 installations.[2] The availability of these OpenPower servers with the Cluster 1350 allows clients to have a homogeneous POWER™ processor-based cluster with binary compatibility for all software versions across all nodes.

And, for clients who prefer a standard rack-mount server for their native 64-bit computing platform the Cluster 1350 now offers two new POWER5 processor-based servers. The IBM System p5 505 provides the performance and reliability of the POWER5 architecture in a space-saving 1U package. Clients will find this platform ideal for HPC Linux cluster workloads that demand superior floating point capabilities as well as scale out and infrastructure workloads such as file/print, web serving, networking, security, and systems management. In addition, the new IBM System p5 550/550Q server is an excellent node for compute-intensive business intelligence and data warehouse applications, as well as general HPC applications that require large local storage and high I/O bandwidth.

Reliable, tested configurations

Standard configurations of the Cluster 1350 include compute nodes, at least one management node (one redundant management node for failover is optional), and up to 64 storage nodes for a maximum of 1,024 managed nodes. Clients who require configurations larger than a total of 1,026 nodes (1,024 managed nodes and up to two management nodes) or components not included in standard Cluster 1350 configurations can use a special bid process to request support of these custom configurations. All of these larger configurations utilize standard 42U racks.

Smaller cluster environments may use the Cluster 1350 25U racks, which allow clients to optimize the size and affordability of their cluster to meet specific application needs. For example, clients with database, business intelligence (BI) and general SMB (small and medium business) applications will find that these smaller racks enable extremely cost-effective solutions for the smaller cluster configurations normally required in these environments.

Cluster 1350 clients may choose from a broad variety of cluster interconnect technologies from several of the industry's leading network switch and adapter vendors to meet the specific performance needs of their cluster application environment. These choices span the full range of high-performance networking technologies including Gigabit Ethernet, InfiniBand and Myrinet switches and adapters. In addition, each Cluster 1350 includes a management Ethernet VLAN for highly secure internode communication.

Expanding possibilities

The Cluster 1350 also offers clients the opportunity to take advantage of GPFS for Linux to expand and enhance their high-performance cluster data storage requirements. GPFS is a high-performance, scalable, shared-disk file system that provides fast data access from all nodes in a Linux cluster and NFS export capabilities outside the cluster. Parallel applications running across multiple nodes of the cluster as well as serial applications running on a single node can readily access shared files using standard UNIX file system interfaces. Furthermore, GPFS can be configured for failover from both HDD and server malfunctions.

In addition, the Cluster 1350 incorporates the technology of the IBM TotalStorage® DS4000 to provide highly reliable data storage for business-critical applications that require high-speed transfer and large amounts of data. Optional cluster components include IBM TotalStorage, DS4100, DS4300, DS4300 Turbo and DS4500 Storage Servers along with DS4000 EXP100 and the EXP400 Ultra 320 SCSI Storage Expansion Units. In addition, the TotalStorage SAN 16B-2 SAN switch is now available to provide a robust SAN solution for cluster applications.

Summary

Creating a computing infrastructure is an exercise in balancing price and performance to deliver the appropriate solution for each client's specific business needs.

For many high-performance workloads, the most advantageous solution is clustering. Harnessing the power of multiple servers in parallel allows for the management and resolution of computationally intense problems with an excellent price/performance ratio. Clustering is also an excellent approach for consolidating multiple workloads, thereby enhancing manageability and availability. In addition, the advent of Linux as a platform for building powerful clustered systems offers clients access to the growing knowledge base and expert contributions of the open source community.

The IBM @server Cluster 1350 is a comprehensive solution that can help simplify and expedite deployment of a Linux cluster. IBM combines all hardware, software and services into a single product offering, allowing clients the benefit of a single point-of-contact for the entire cluster rather than dealing with multiple vendors for individual components.

The Cluster 1350 is the solution of choice for any organization that recognizes the economic advantages of deploying a Linux cluster, but has concerns about the time and technical resources required for the end-to-end implementation.

IBM @server Cluster 1350 at a glance

Building block	xSeries 346	xSeries 336	@server 326m	BladeCenter HS20	BladeCenter JS20
Node type	Management/ compute/storage	Compute/storage	Management/ compute/storage	Compute/storage	Compute/storage
Packaging	Rack drawer (2U)	Rack drawer (1U)	Rack drawer (1U)	14 per rack drawer (7U)	14 per rack drawer (7U)
Processor	3.0, 3.2, 3.4, 3.6, 3.8 GHz Intel Xeon	2.8, 3.0, 3.2, 3.4, 3.6, 3.8 GHz Intel Xeon	2.6, 2.8 GHz (single-core) AMD Opteron; 1.8, 2.0, 2.2, 2.4 GHz (dual-core) AMD Opteron	3.2, 3.4, 3.6, 3.8 GHz Intel Xeon	2.2 GHz PowerPC 970
L2 cache	1MB	1MB	1MB	1MB	512KB
Memory (standard)	512MB or 1GB	512MB or 1GB	1GB	512MB or 1GB	512MB
HDD bays	6 (hot-swap SCSI)	2 (fixed SATA or hot-swap SCSI) or 4 hot-swap SFF* SCSI	2 (fixed SATA or hot-swap SCSI)	2 (fixed SCSI or SFF* fixed SCSI)	2 (fixed IDE)
I/O expansion slots	4 PCI-X (64-bit) or 2 PCI-X and 2 PCI-E	2 PCI-X (64-bit) or 1 PCI-X and 1 PCI-E	2 PCI-X (64-bit)	1 PCI-X (64-bit)	1 PCI-X (64-bit)
Integrated storage controllers	Ultra320 SCSI	Ultra320 SCSI or SATA	Ultra320 SCSI or SATA	Ultra320 SCSI	IDE
Integrated Ethernet controllers	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet
System expansion					
Max memory	16GB	16GB	16GB	8GB	8GB
SCSI internal storage (minimum/maximum)	36GB/1.2TB	36GB/600GB	36GB/600GB	36GB/146GB	—
IDE internal storage (minimum/maximum)**	—	—	—	—	40GB/80GB
SATA internal storage (minimum/maximum)	—	80GB/500GB	80GB/160GB	—	—
System connectivity for Cluster VLAN	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Myrinet, InfiniBand

* Small form factor

** Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available.

Building block	@server AMD Opteron LS20	OpenPower 710[2]	OpenPower 720[2]	p5-505	p5-550/550Q
Node type	Compute	Management/ storage	Management/ storage	Compute/ management/ storage	Compute/ management/ storage
Packaging	14 per rack drawer (7U)	Rack drawer (2U)	Rack drawer (4U)	Rack drawer (1U)	Deskside or rack drawer (4U)
Processor	2.0, 2.4, 2.6 GHz single-core AMD Opteron; 2.0, 2.2 GHz dual-core AMD Opteron	1.65 GHz POWER5	1.5, 1.65 GHz POWER5	1.5, 1.65 GHz POWER5	1.5, 1.9 GHz POWER5+
L2 cache	1MB—single-core 1MB/core—dual-core	1.9MB	1.9MB	1.9MB	1.9MB per processor pair
Memory (standard)	1GB or 2GB	512MB or 1GB	512MB or 1GB	1GB to 32GB	1GB to 64GB
HDD bays	2 (SFF* fixed SCSI)	4 (hot-swap SCSI)	Four standard, plus four optional (10K or 15K rpm disks)	Two (10K or 15K rpm disks)	Four standard, plus four optional (10K or 15K rpm disks)
I/O expansion slots	1 PCI-X (64-bit)	3 PCI-X (64-bit)	5 hot-plug PCI-X (64-bit)	PCI-X 2.0 (two 266 MHz (DDR))	5 PCI-X (4 – 133 MHz; 1 – 266 MHz (DDR))
Integrated storage controllers	LSI 1020 SCSI	Ultra320 SCSI (dual-ported)	Ultra320 SCSI (dual-ported)	Ultra320 SCSI	Ultra320 SCSI
Integrated Ethernet controllers	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet	Dual Gigabit Ethernet
System expansion					
Max memory	8GB	32GB	64GB		
SCSI internal storage (minimum/maximum)	36GB/146GB	73.4GB/1.2TB	73.4GB/2.4TB	73.4GB/600GB	73.4GB/2.4TB
IDE internal storage (minimum/maximum)**	—	—	—	—	—
SATA internal storage (minimum/maximum)	—	—	—	—	—
System connectivity for Cluster VLAN	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Other	Gigabit Ethernet, Other	Gigabit Ethernet, Myrinet, InfiniBand	Gigabit Ethernet, Myrinet, InfiniBand

* Small form factor

** Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available.

IBM @server Cluster 1350 at a glance (cont.)

External storage

Storage servers	
1724-10U	IBM TotalStorage DS4100 Storage Server (formerly FAStT100)
1722-60U	IBM TotalStorage DS4300 Storage Server (formerly FAStT600)
24P8215	IBM TotalStorage DS4300 Turbo (formerly FAStT600 with Turbo option)
1742-90U	IBM TotalStorage DS4500 Storage Server (formerly FAStT900)
Storage expansion	
1710-10U	IBM DS4000 EXP100 Storage Expansion Unit (formerly FAStT EXP100)
1733-1RU	IBM EXP400 Ultra 320 SCSI Storage Expansion Unit
SAN switch	
2005-16B	IBM TotalStorage SAN 16B-2 SAN Switch

Software

Operating system[1] (Linux distributions)	• Red Hat Enterprise Linux 4 (RHEL 4) • SUSE Linux Enterprise Server 9 (SLES 9)
Cluster management software	• IBM Cluster Systems Management (CSM) for Linux V1.5 (optional) • IBM General Parallel File System (GPFS) for Linux V2.3 (optional) • Scali MPI Connect (optional)
Rack dimensions	
42U primary rack/expansion rack	79.5" H x 25.2" W x 43.3" D (2020 mm x 640 mm x 1100 mm); 574.2 lbs (261 kg)
25U rack	49.0" H x 23.8" W x 39.4" D (1344 mm x 605 mm x 1001 mm); 221 lbs (100.2 kg)
Scalability	One management node is required and one redundant management node for failover is optional. A minimum of two and a maximum of 1,024 managed nodes are supported (optional software may limit supported nodes). This total can include up to 64 storage nodes. Therefore, the minimum configuration includes three nodes (one management node and two compute nodes). The maximum configuration is 1,026 nodes including compute, storage and management nodes. Larger configurations are available through a special bid process.
Services	• Cluster 1350 hardware installation is included at no charge on 42U and 25U racks. • Linux cluster software services and SupportLine for Linux Clusters are available as optional fee-based services.
Warranty	Basic limited warranty service on most IBM and non-IBM components: three year, next business day, on-site support for selected components. Some xSeries 336 and @server 326m models have a one year warranty. Enhanced warranty service plans are available.

For more information

To learn more about the IBM @server Cluster 1350, contact your IBM representative or IBM Business Partner, or visit the following Web sites:

- **ibm.com**/eserver/clusters
- **ibm.com**/common/ssi
- www.redhat.com/software/rhel
- www.suse.com/us/business/index.html

Information concerning non-IBM products was obtained from the suppliers of these products. Questions on the capabilities of the non-IBM products should be addressed with the suppliers. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven® or ClusterProven®.

When referring to storage capacity, 1TB equals total GB divided by 1000; accessible capacity may be less.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

[1] Clients may acquire a supported version of the Linux operating system specified by IBM (including optional drop-in-the-box offerings) from a Linux distributor or an authorized IBM Business Partner.

[2] OpenPower 710 and 720 servers are available with the Cluster 1350 via a special bid process.

© Copyright IBM Corporation 2006

IBM Corporation
Integrated Marketing Communications,
Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States of America
January 2006
All Rights Reserved

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features or services discussed in this publication in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the e-business logo, @server, AIX 5L, BladeCenter, ClusterProven, OpenPower, POWER, POWER5, POWER5+, Power PC, pSeries, ServerProven, System p5, X-Architecture, xSeries and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A full list of U.S. trademarks owned by IBM may be found at **ibm.com**/legal/copytrade.shtml.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

CLD00221-USEN-10
GM13-0159-10